

FTB events

In this document we present the list of FTB events used by MVAPICH2 to proactive migrate MPI processes when a node is failing.

CR_FTB_NODE_FAIL

This event is published by the Migration-Trigger component and consumed by mpirun_rsh (the job launcher). It contains the name of the failing node.

Severity: INFO

Payload: node:<hostname> The payload includes the hostname which is expected to fail.

CR_FTB_USER_TRIGGER

This event is published when a user requests a migration (for example in a maintenance operation). Published by the Migration-Trigger component and consumed by the job launcher. It contains the name of the source node of migration.

Severity: INFO

Payload: node:<hostname> The payload includes the hostname which is expected to become unavailable.

CR_FTB_MIGRATE

This event is published by mpirun_rsh and consumed by the mpispawns and CR (Checkpoint Restart) module. It contains the name of the source and target nodes of the migration.

Severity: INFO

Payload: node:<src_hostname, dest_hostname> The payload contains the hostname of the source and target nodes of the migration.

CR_FTB_CKPT_DONE

This event is published by the MPI processes that were able to checkpoint and indicates that the checkpoint is completed successfully. It is consumed by mpirun_rsh.

Severity: INFO

CR_FTB_CKPT_FAIL

This event is published by the MPI processes that failed to take the checkpoint and indicates that the checkpoint is failed. It is consumed by mpirun_rsh.

Severity: ERROR

CR_FTB_MIGRATE_IMG_C

This event is published by the mpispawn on the migration source and indicates that the migration of the checkpoint image is complete. It is consumed by mpirun_rsh.

Severity: INFO

Payload: node:<src_hostname, tgt_hostname> The payload includes the source and target nodes of the migration.

CR_FTB_MIGRATE_RESTART

This event is published by mpirun_rsh and consumed by mpispawns. It indicates that the migrated processes can be restarted.

Severity: INFO

Payload: node:<tgt_hostname>, rank:<mpiranks> The payload includes the target node of the migration and the list of processes ranks migrated.

CR_FTB_RSRT_DONE

This event is published by MPI processes and consumed by mpirun_rsh. It indicates that the restart completed successfully.

Severity: INFO

CR_FTB_RSRT_FAIL

This event is published by MPI processes and consumed by mpirun_rsh. It indicates that the restart failed.

Severity: ERROR

CR_FTB_CKPT_FINALIZE

This event indicates that the CR module has been shutdown. Published by all MPI processes doing MPI_Finalize().

Severity: INFO